

HOTCHIPS26

SX-ACE Processor: NEC's Brand-New Vector Processor

Shintaro Momose, Ph.D.

NEC Corporation

IT Platform Division

Manager of SX vector supercomputer development

August 11th, 2014

SX History and Technical Evolutions

NEC has always provided the high sustained performance by Vector Super-Computer SX series.

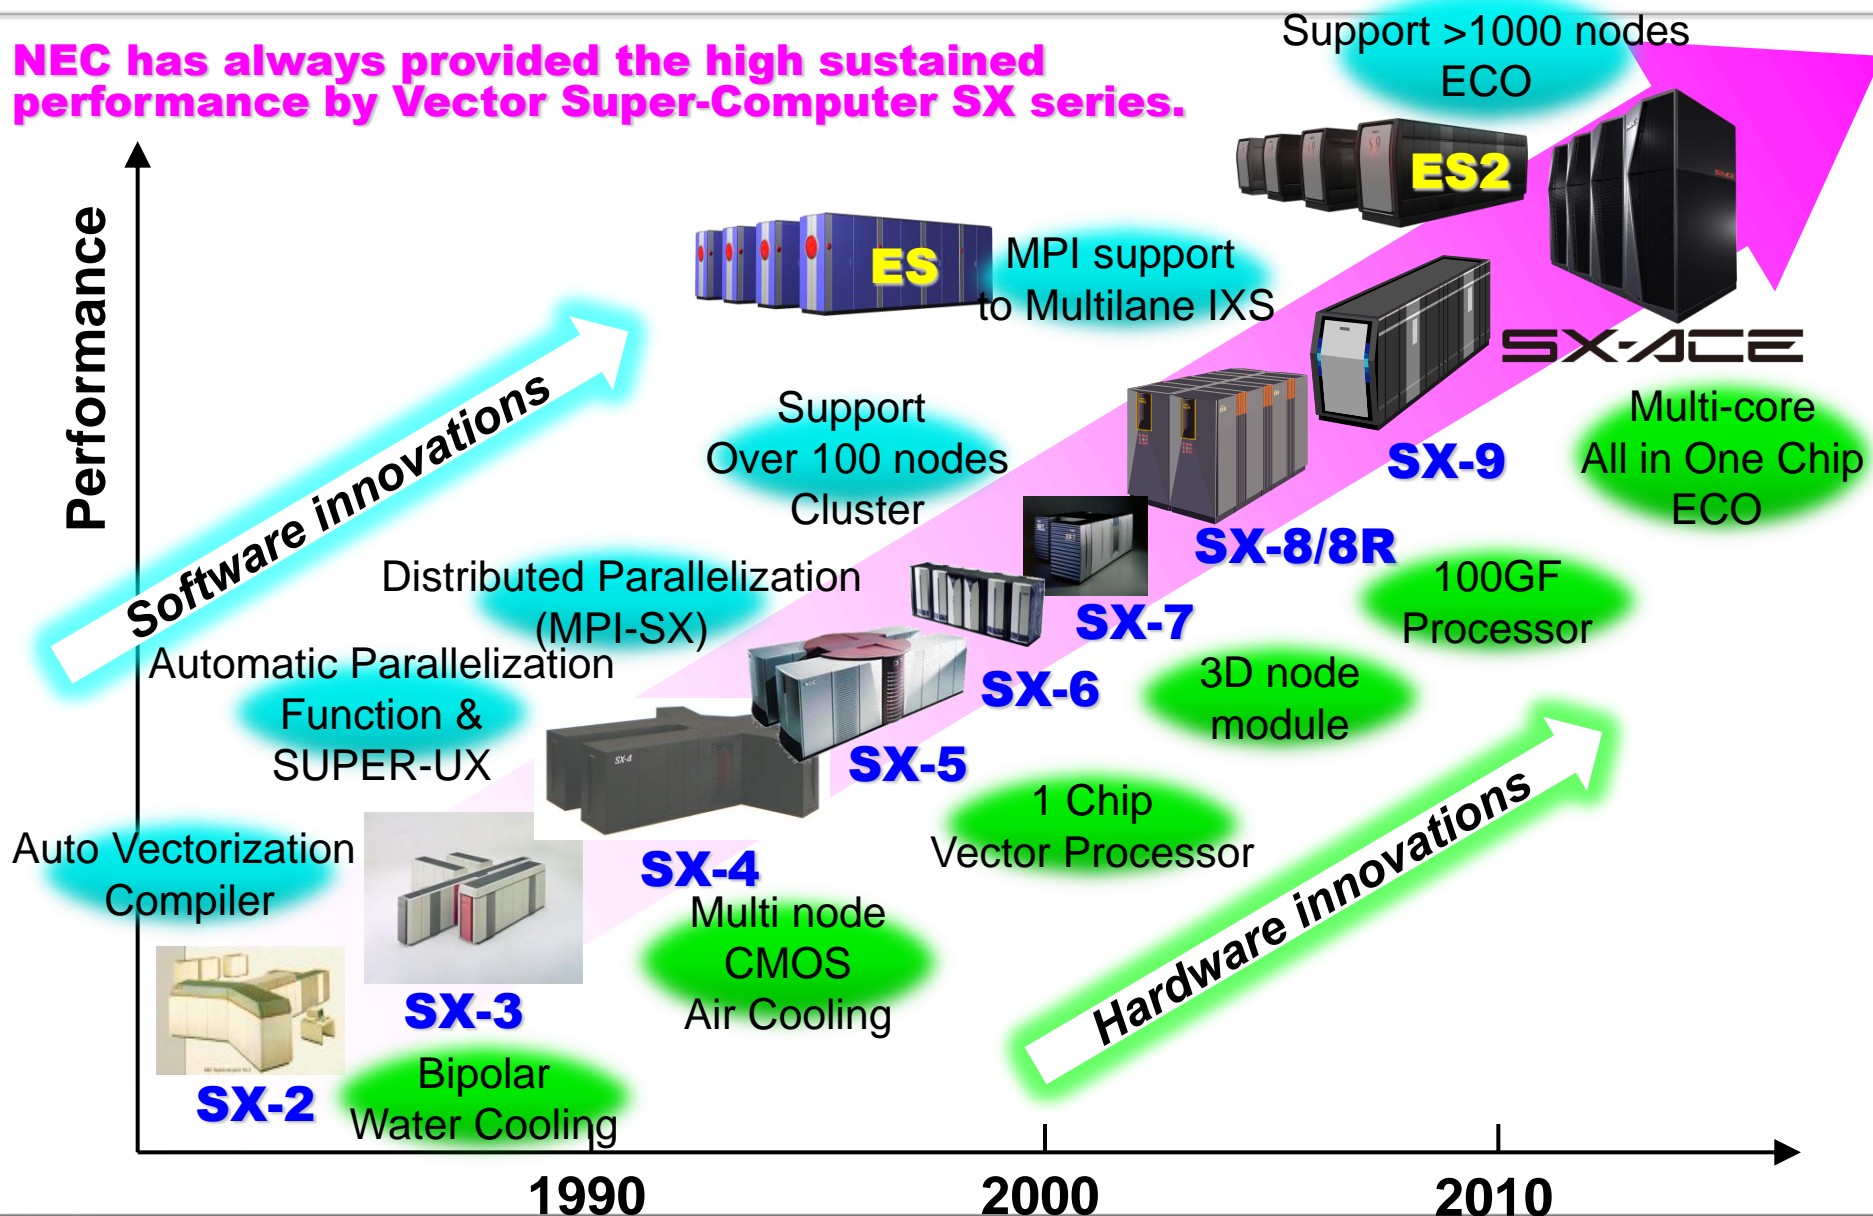


Table of Contents

■ Introduction

- HPC Technical Trends and Issues
- SX-ACE Processor Design Direction

■ Processor Architecture

- Processor/Core Architectures and Implementations
- Memory Subsystem

■ Performance Evaluations


- Experimental Results of Several memory intensive Benchmarks

■ Conclusions

Introduction

Trend of TOP500 (1st ~ 10th system)

- Growing of LINPAC performance has been provided by system enlarging
- User must spend their time to extract massive parallelism
- Smaller # of cores with big cores can reduce the difficulty

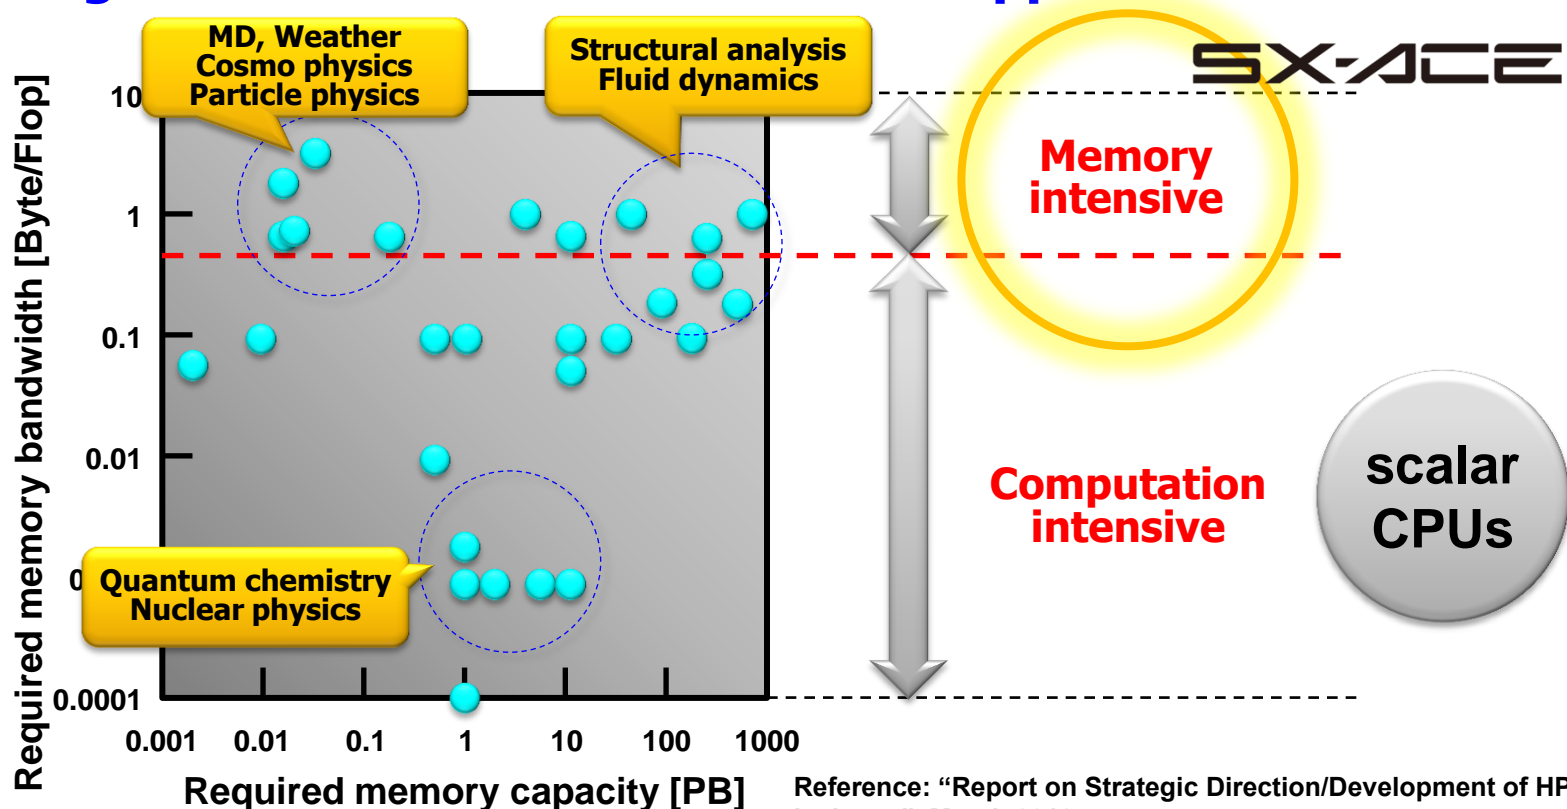


Required Byte/Flop in Real Applications

According to Japanese Government (MEXT) working group report for a wide variety of strategic segment applications, diverse characteristics are observed.

MEXT: Ministry of Education, Culture, Sports, Science & Technology

**B/F requirement from each application differs greatly.
Any single architecture cannot cover all application areas.**



Concepts of SX-ACE



**The best solution for memory intensive APs
against scalar processors trend**

Big Core

Reducing Massive Parallel Difficulty with fewer cores

World top-level performance: 64GF
The largest memory bandwidth: 64~256GB/s

Low Power Consumption

The best memory bandwidth solution

GB/s / Watt 1.8x
compared x86 CPU

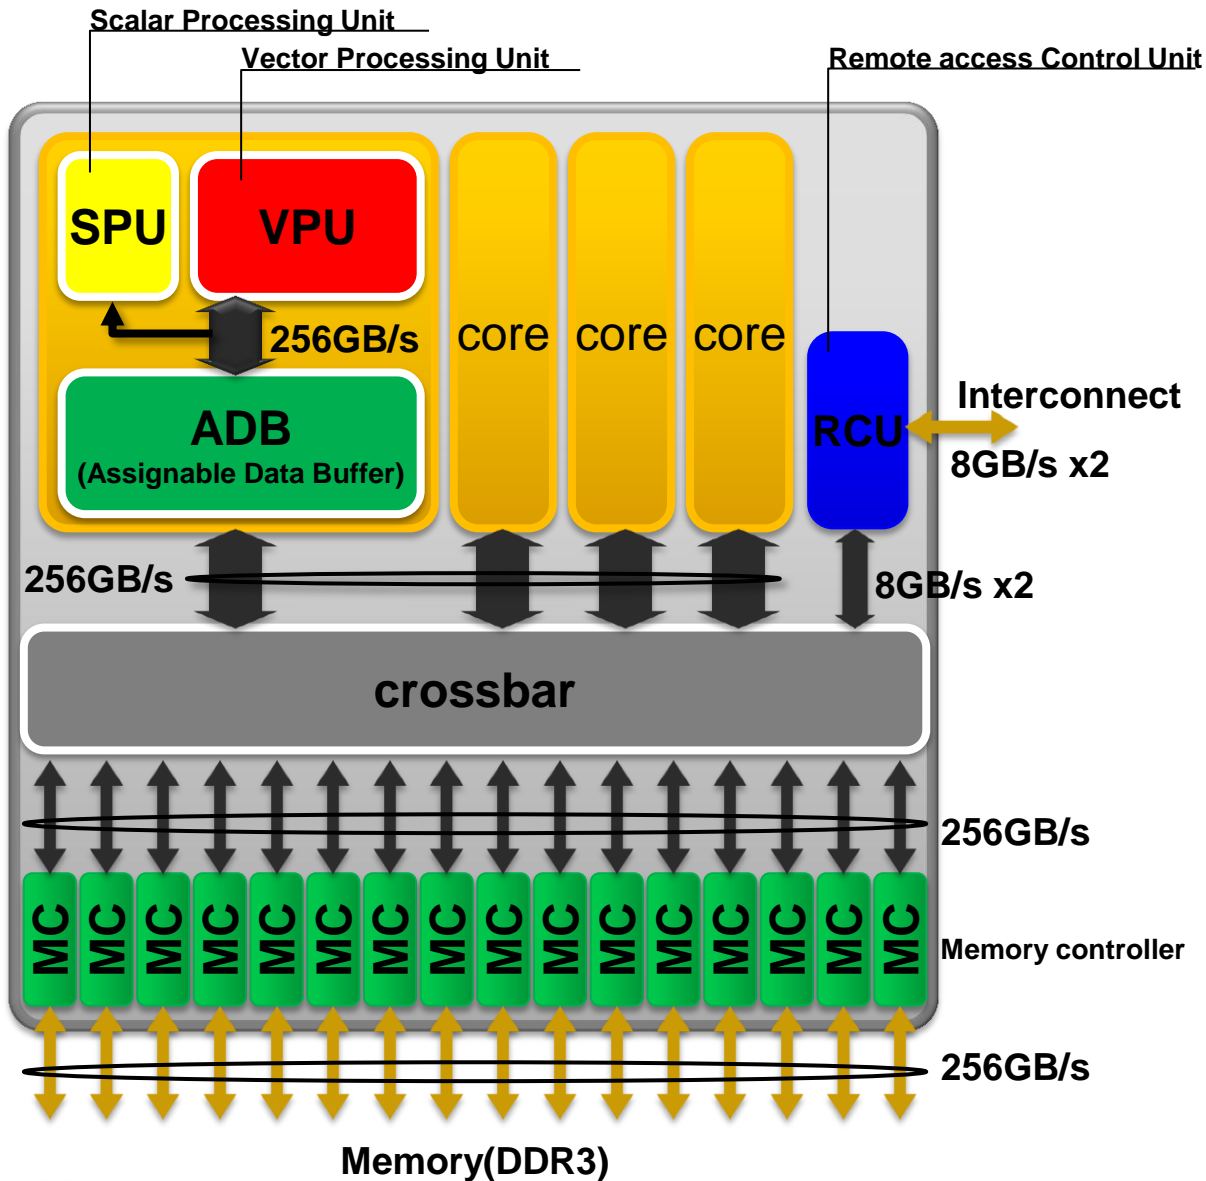
Hybrid Solution

Vector / Scalar tightly coupled environment

Specialized SWs

Architecture


Processor Overview



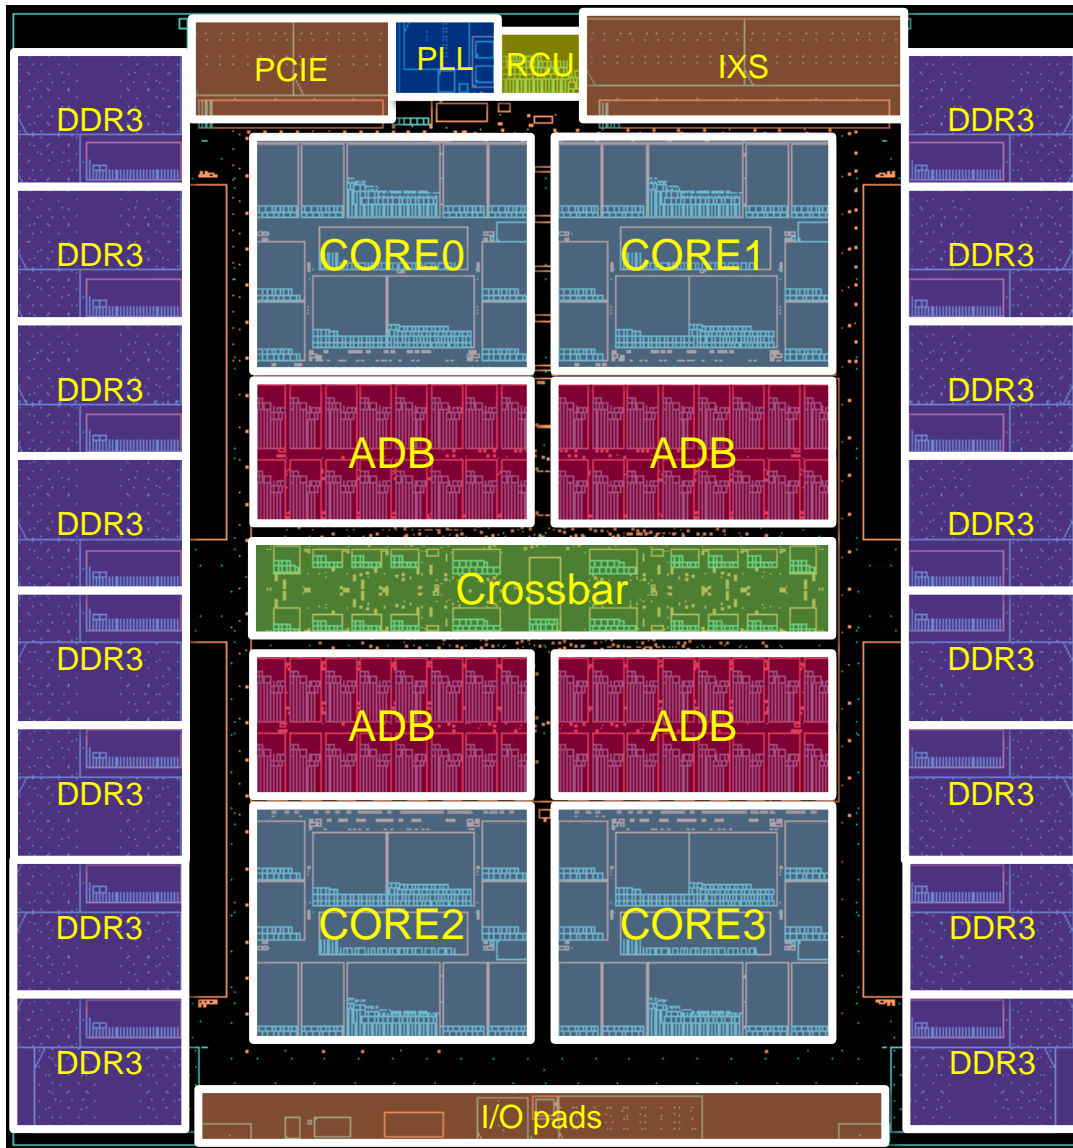
CORE	
Architecture	Vector
Clock Frequency	1.0GHz
SPU decode rate	4 instructions
VPU Performance	64GFlops
ADB size	1MB
ADB bandwidth	256GB/s
Memory bandwidth	64GB/s~256GB/s
Core Byte/Flop	1.0 ~ 4.0
CPU	
Cores	4
Performance	256GFlops
Memory bandwidth	256GB/s
CPU Byte/Flop	1.0
Memory capacity	64GB

Single Core Comparison

The SX-ACE core can provide the world top-level performance and the largest memory bandwidth



Floor Plan of the CPU



■ “Memory access” focused layout

■ Specifications

- Process rule: 28nm
- Clock speed: 1GHz
- Die size:
23.05 x 24.75mm
- # of transistors: 2BTr.

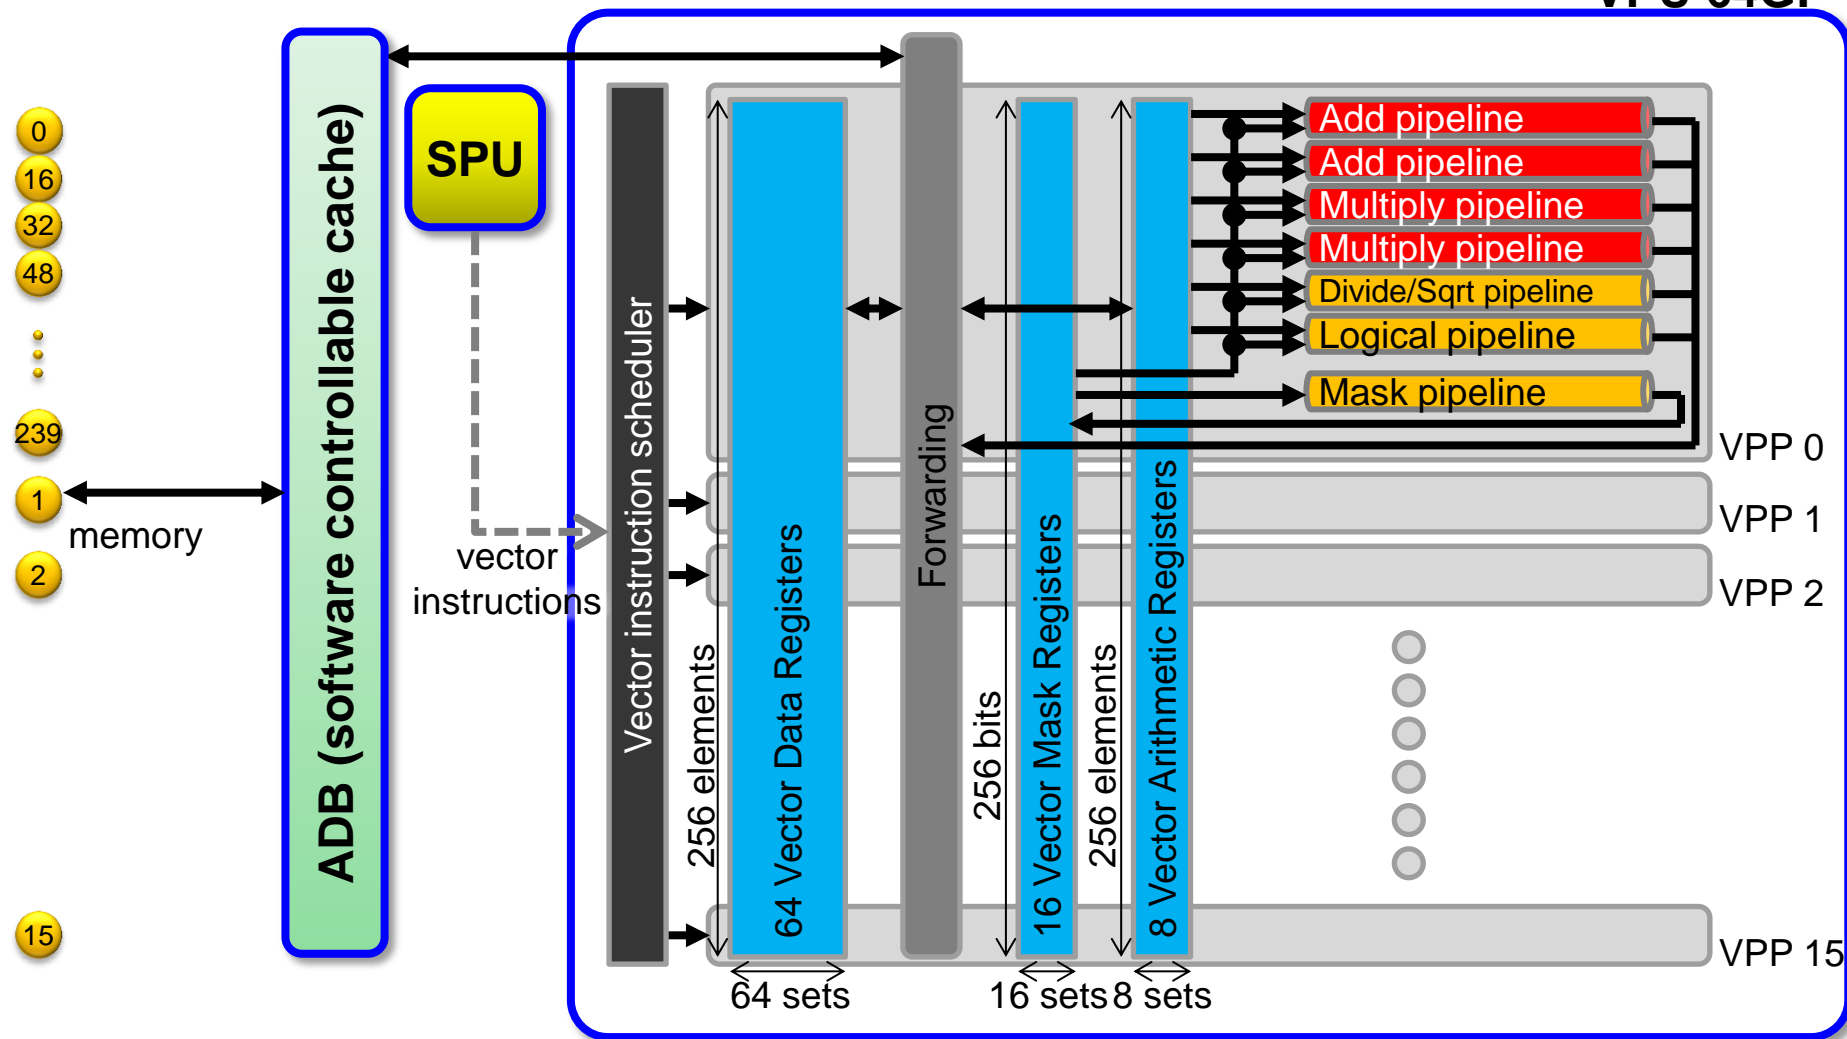
■ I/F

- 16ch DDR3 I/F
- IXS 8GB/s x 2
- 2ch PCIEx8 I/F

Core Architecture

256 operations = 16 parallel x 16 clock cycles

VPU 64GF

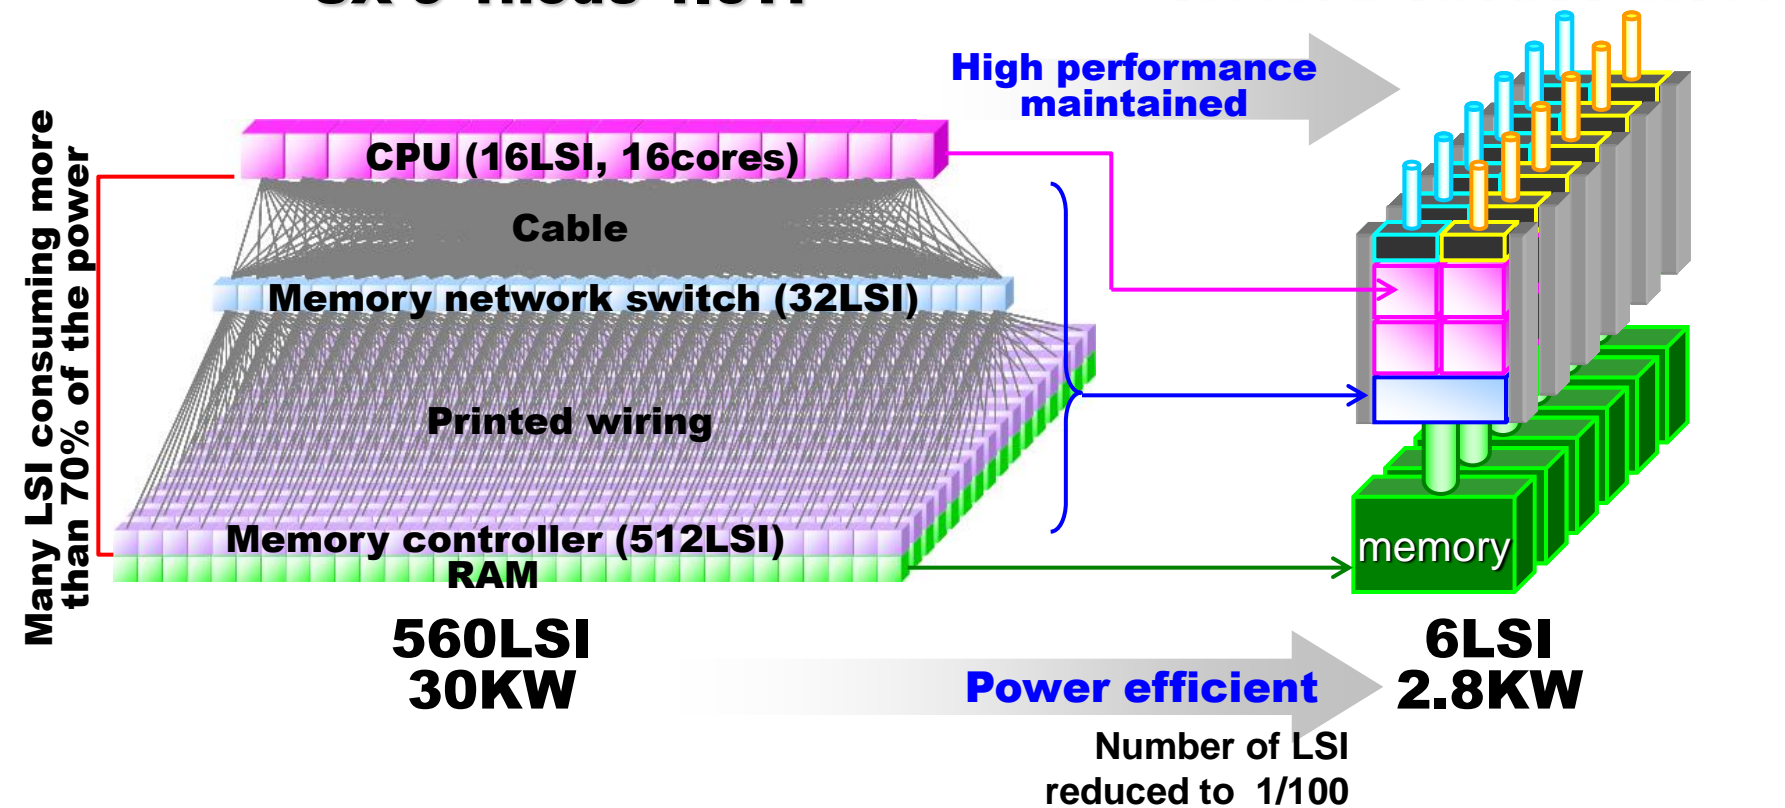


Memory Network Integration

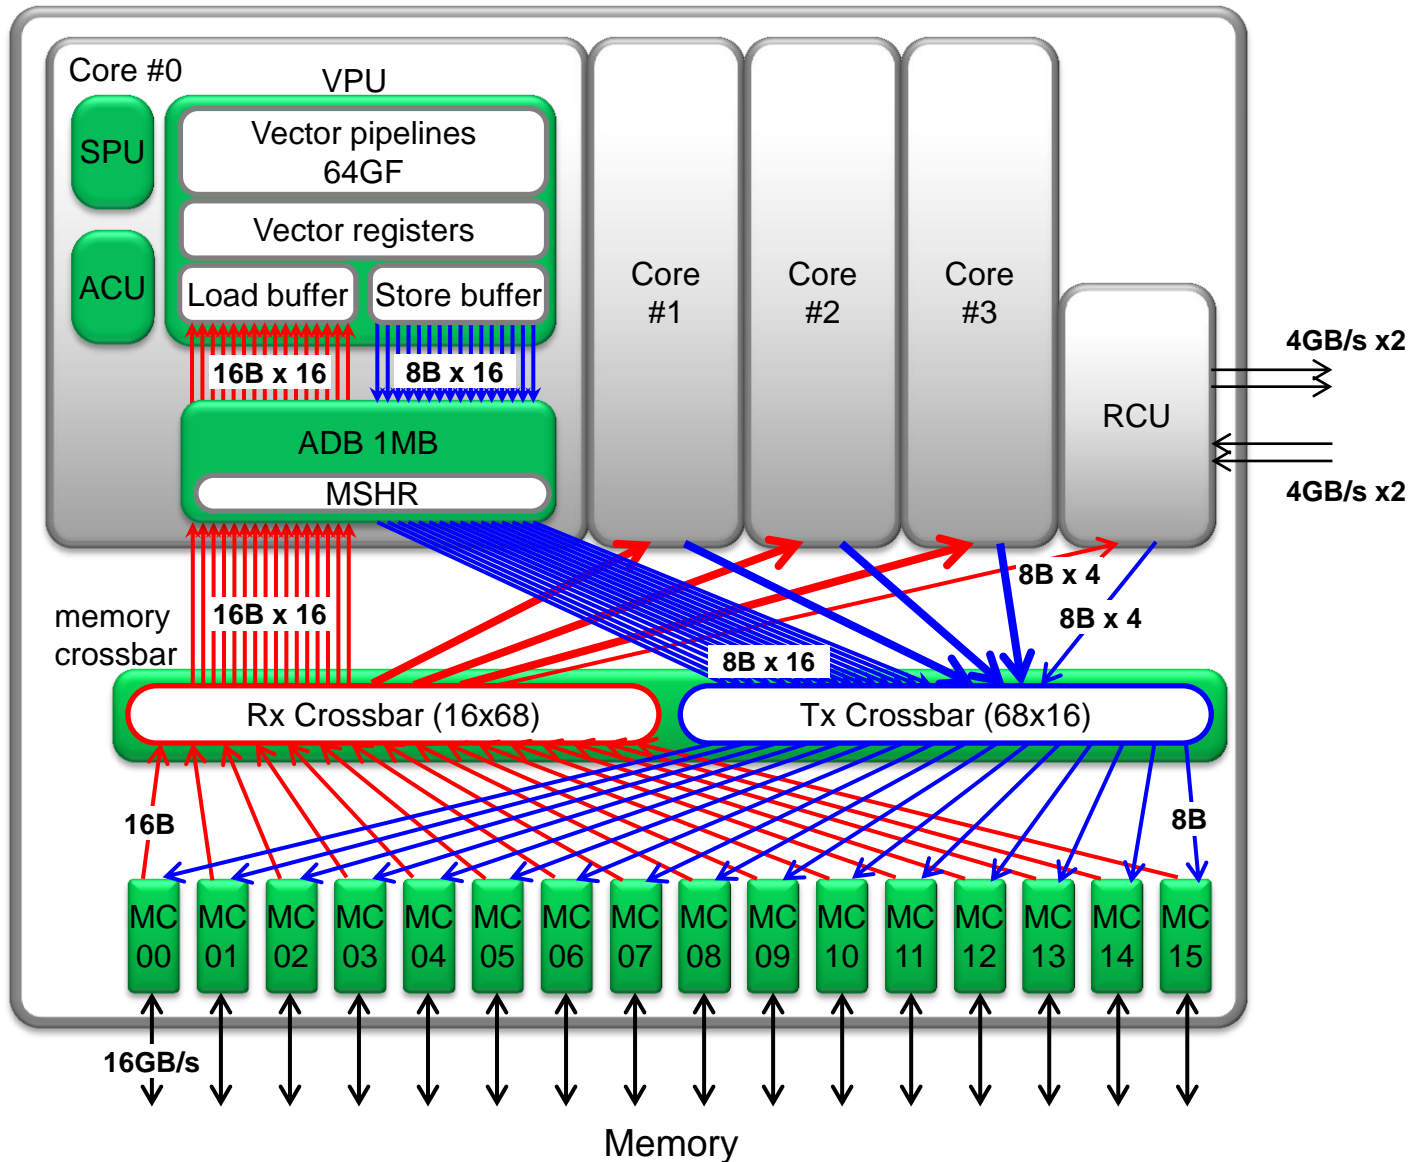
- Large SMP configuration can provide high sustained performance
- But, over 70% power was consumed by the memory network
- **SX-ACE processor integrates the memory network into LSI**

SX-9 1node 1.6TF

SX-ACE 6nodes 1.5TF



Memory Subsystem




Reducing DRAM Energy

■ Cache line size

- DRAM activation powers are depending on
- Sustained memory bandwidth is strongly affected by adopted cache line size

■ Variable cache line size feature

- Supporting 64B/128B memory access granularity
- 128B as a default to reduce power
- 64B for a sparse memory access such as stride/indirect memory accesses



RD:WR 1:1, Micron DDR3 power calculator 0.96

Assignable Data Buffer (ADB)

■ On-chip Cache for Vector


- Private, 1MB, 4-way, 16-bank
- 256GB/s bandwidth per core
- Software controllable cache
- Customized for fast random access

■ Assignable Feature

- A bypass flag in each instruction
- Compiler/User can control
- Avoiding cache pollution


■ MSHR Feature

- Redundant memory requests same as an inflight memory request are held to reduce memory transactions




Out-of-Order Vector Memory Access


Vector memory access instruction



Consecutive memory access (by HW)



Stride memory access (by HW)




Indirect memory access (by SW)



Node Packaging

CPU

256GFlops = 64GFlops/core x 4
256GB/s memory bandwidth




Memory

16 DIMMs (DDR3 2000)
256GB/s, 64GB

Rated power consumption = 469W

Hybrid Cooling



Optimization of cooling efficiency and rack weight

- CPU: water cooling
- Other components: air cooling

Performance Evaluation

Performance Evaluation Conditions

Evaluation programs

Evaluate point	Benchmark
Off-chip memory bandwidth	STREAM (TRIAD)
Off/On-chip memory bandwidth	Himeno Benchmark (High memory intensive)
Indirect memory access performance	Legendre transformation

Each evaluation is carried out by only using compiler optimizations without code modifications for individual systems

Performance comparison

CPU	Performance	Memory bandwidth	Rated system Watts/CPU
SX-9	102GF = 102GF x 1c	256GB/s	1875W
SX-ACE	256GF = 64GF x 4c	256GB/s	469W
IVB(Xeon)	230GF = 19GF x12c	60GB/s	200W
Power7	245GF = 31GF x 8c	128GB/s	656W
FX10(Sparc)	234GF = 15GF x16c	85GB/s	281W


Power7 and FX10 are measured through a joint research with Tohoku University

Memory Bandwidth 1


Evaluation of Off-chip memory bandwidth

■ Benchmark code: STREAM (TRIAD)

Sustained memory bandwidth



Power efficiency (SX-ACE=1)



- Only the SX-ACE single core can use full memory bandwidth
- This can accelerate memory-intensive serial parts in parallel processing


- SX-ACE provides the best memory bandwidth per watt

Memory Bandwidth 2


Evaluation of Off/On-chip memory bandwidth

- Benchmark code: Himeno benchmark (highly memory intensive)
solving the Poisson equation with the Jacobi iterative method

Sustained performance (SX-ACE=1)



Power efficiency (SX-ACE=1)



- ADB and MSHR improve sustained memory bandwidth compared with its predecessor
- SX-ACE is the best


- SX-ACE is assumed to provide **2~25x higher power efficiency** in the case of memory intensive APs having off/on chip memory accesses

Indirect Memory Access


Evaluation of Indirect memory access performance

- Benchmark code: Legendre transformation
- Cache effective BM (4.4MB data)

Sustained performance (SX-ACE=1)



Power efficiency (SX-ACE=1)



- Cache is effective
- ADB, MSHR, OoO, and short memory access latency work well

- SX-ACE improvement provides 25x higher power efficiency than SX-9
- But, IVB is the best due to a larger cache and a lower power consumption

Conclusions

■ Issue of modern scalar/accelerator processors

- Massive parallel with small cores
- Low memory bandwidth

■ SX-ACE direction

- Providing the big core with large memory bandwidth
- Improving proven vector architecture

■ SX-ACE processor

- 4 cores vector processor
- 64GF core performance with 64-256GB/s memory bandwidth
- Efficient memory subsystem for higher sustained memory bandwidth

■ Performance

- High sustained performance and power efficiency for memory intensive benchmarks

Acknowledgement

I would like to express my gratitude to Cyber Science Center at Tohoku University for the intensive performance evaluation of the SX vector supercomputers as a part of the joint research project with NEC Corporation.

Tohoku University, Cyber Science Center

- | | |
|-----------------------|-------------------|
| ■ Professor | Hiroaki Kobayashi |
| ■ Associate Professor | Ryusuke Egawa |
| ■ Assistant Professor | Kazuhiko Komatsu |

NEC Corporation

- | | |
|------------------|------------------|
| ■ Senior Manager | Takashi Hagiwara |
| ■ Manager | Yoko Isobe |

Empowered by Innovation

NEC